



US009417971B2

(12) **United States Patent**  
**Agombar et al.**

(10) **Patent No.:** **US 9,417,971 B2**  
(45) **Date of Patent:** **\*Aug. 16, 2016**

(54) **CASCADE ORDERING**

(71) Applicant: **INTERNATIONAL BUSINESS MACHINES CORPORATION**,  
Armonk, NY (US)

(72) Inventors: **John P. Agombar**, Winchester (GB);  
**Christopher B. Beeken**, Winchester (GB); **David J. Carr**, Winchester (GB)

(73) Assignee: **INTERNATIONAL BUSINESS MACHINES CORPORATION**,  
Armonk, NY (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **14/725,312**

(22) Filed: **May 29, 2015**

(65) **Prior Publication Data**

US 2015/0261618 A1 Sep. 17, 2015

**Related U.S. Application Data**

(63) Continuation of application No. 14/600,966, filed on Jan. 20, 2015, which is a continuation of application No. 13/529,691, filed on Jun. 21, 2012, now Pat. No. 9,063,894, which is a continuation of application No. 13/102,886, filed on May 6, 2011, now Pat. No. 8,959,300.

(30) **Foreign Application Priority Data**

May 18, 2010 (EP) ..... 10163131

(51) **Int. Cl.**  
**G06F 11/14** (2006.01)  
**G06F 3/06** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 11/1456** (2013.01); **G06F 3/065**

(2013.01); **G06F 3/0619** (2013.01); **G06F 3/0655** (2013.01); **G06F 3/0665** (2013.01); **G06F 3/0688** (2013.01); **G06F 3/0689** (2013.01); **G06F 11/1451** (2013.01); **G06F 11/1458** (2013.01); **G06F 11/1464** (2013.01);  
(Continued)

(58) **Field of Classification Search**

None

See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

4,181,933 A 1/1980 Benysek  
4,974,156 A 11/1990 Harding et al.

(Continued)

**OTHER PUBLICATIONS**

Definition of "metadata", Aug. 9, 2013, 1 pageFree Online Dictionary of Computing.

(Continued)

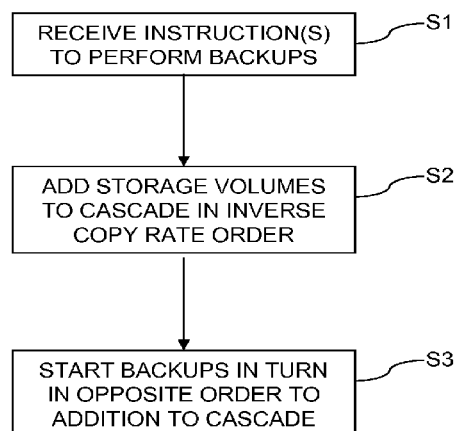
*Primary Examiner* — Yaima Rigol

(74) *Attorney, Agent, or Firm* — Griffiths & Seaton PLLC

(57) **ABSTRACT**

A method comprises creating a cascade of target storage volumes from a plurality of target storage volumes by inserting each most recent target storage volumes between a last added target storage volume and a source storage volume, the most recent target storage volumes each having a copy rate faster than the last added target storage volume, such that each added target storage volume has a faster copy rate than an immediately prior added target storage volume, and a most recent target storage volume added to the cascade of target storage volumes has a fastest copy rate; and receiving one or more instructions initiating a plurality of backup processes from the source storage volume to a plurality of target storage volumes.

**8 Claims, 6 Drawing Sheets**



(52) U.S. Cl.

CPC ..... **G06F 11/1466** (2013.01); *G06F 11/1461*  
 (2013.01); *G06F 2003/0692* (2013.01); *G06F*  
*2201/82* (2013.01); *G06F 2201/84* (2013.01)

(56)

**References Cited**

## U.S. PATENT DOCUMENTS

5,504,882	A	4/1996	Chai et al.
7,360,048	B2	4/2008	Agombar et al.
7,386,695	B2	6/2008	Fuente
7,676,641	B2	3/2010	Agombar et al.
8,219,770	B2	7/2012	Usami
2006/0080362	A1	4/2006	Wagner et al.
2006/0139697	A1	6/2006	Fuente
2006/0143413	A1	6/2006	Agombar et al.
2006/0218368	A1	9/2006	Satoyama et al.
2007/0233987	A1	10/2007	Maruyama et al.
2007/0283090	A1 *	12/2007	Kaji ..... G06F 11/2094 711/114

2008/0162844	A1	7/2008	Yoder et al.
2008/0306989	A1	12/2008	Boyce et al.
2008/0320260	A1	12/2008	Watanabe
2009/0077338	A1	3/2009	Agombar et al.
2009/0319735	A1	12/2009	Agombar et al.
2009/0327589	A1	12/2009	Moshayedi
2009/0327626	A1	12/2009	Kaushik et al.
2010/0042795	A1	2/2010	Uchida
2010/0169596	A1	7/2010	Wada
2011/0231698	A1	9/2011	Zlati et al.
2011/0283072	A1	11/2011	Uratani et al.
2012/0246424	A1	9/2012	Yuhara et al.

## OTHER PUBLICATIONS

IBM TotalStorage Enterprise Storage Server Implementing ESS  
 Copy Services in Open Environments, 642 pages, IBM.

\* cited by examiner




Fig. 1




Fig. 2

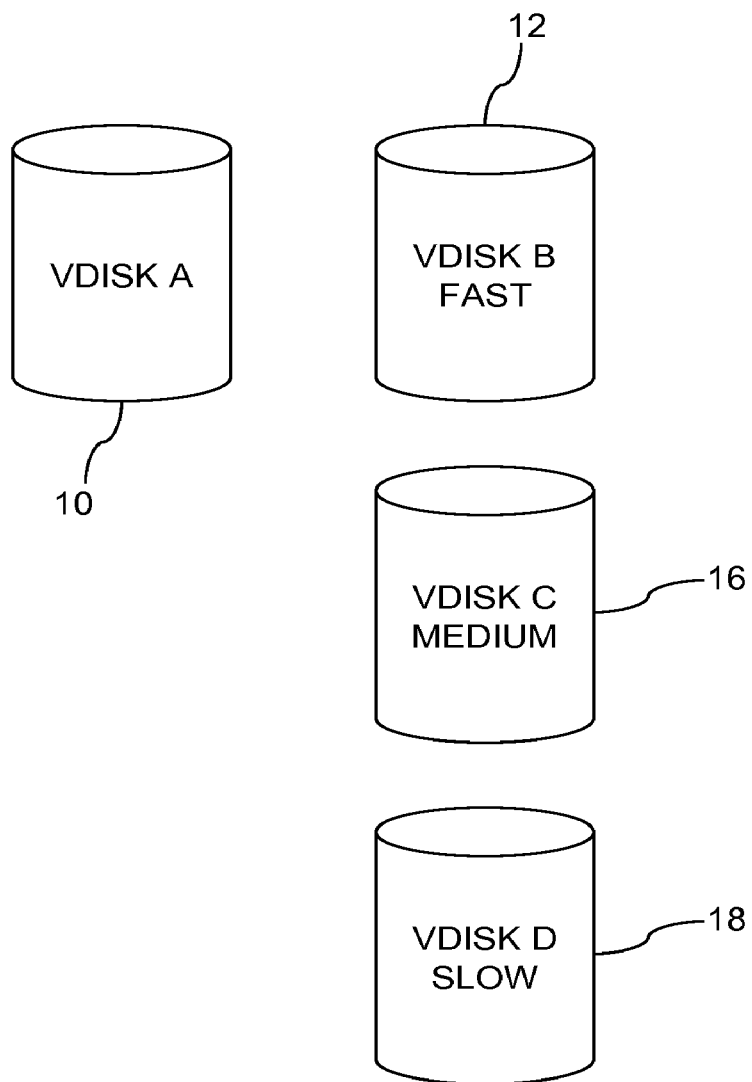


Fig. 3

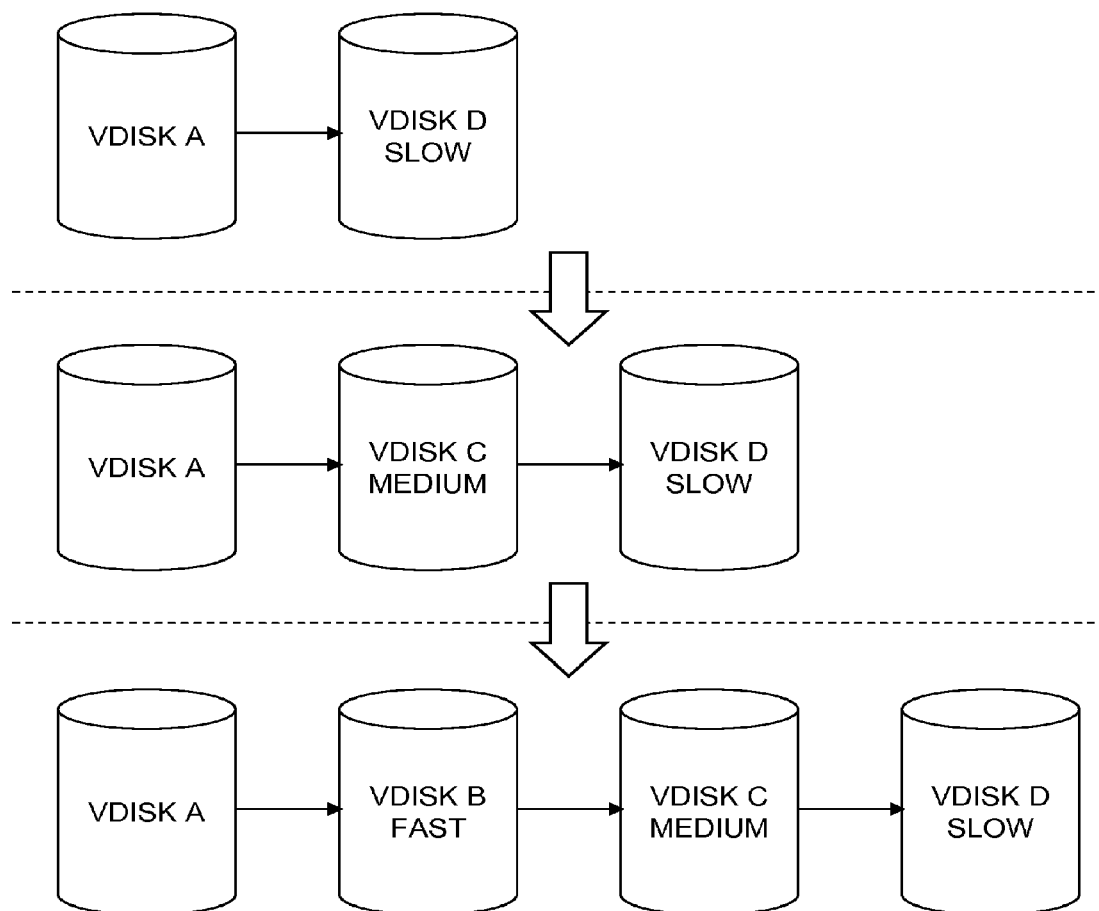


Fig. 4

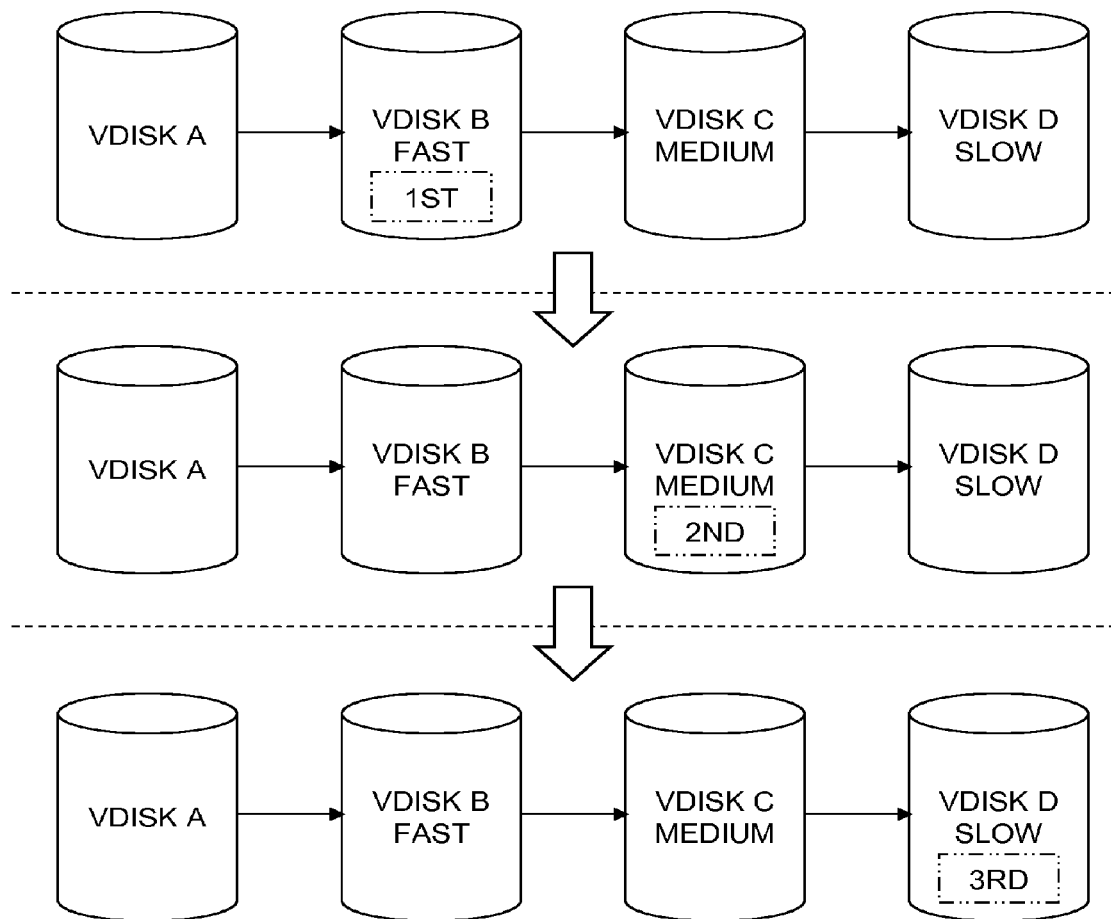


Fig. 5

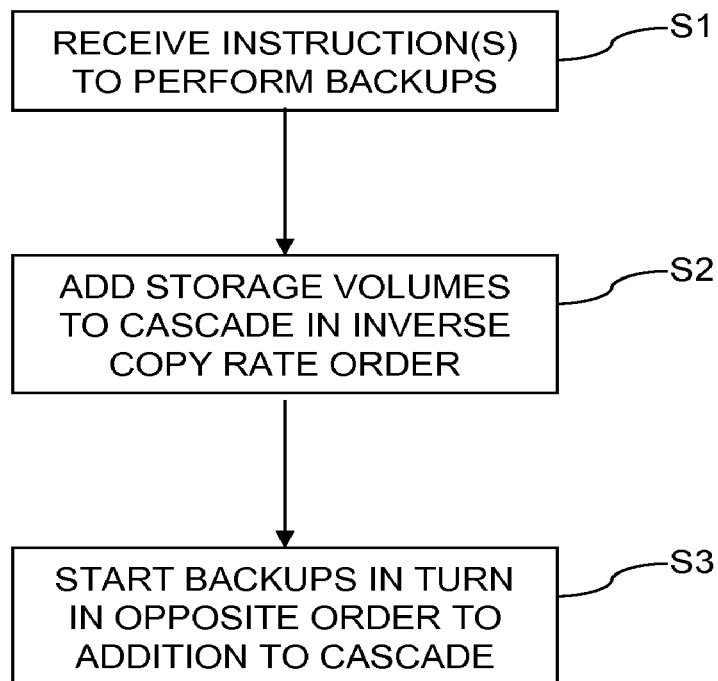


Fig. 6



**CASCADE ORDERING****CROSS REFERENCE TO RELATED APPLICATIONS**

This application is a Continuation of U.S. patent application Ser. No. 13/529,691 filed on Jun. 21, 2012, which is a Continuation of U.S. patent application Ser. No. 13/102,886, filed on May 6, 2011, which claims priority to European Patent Application No. 10163131.5, filed May 18, 2010, both of which are hereby incorporated by reference.

**BACKGROUND OF THE INVENTION****1. Field of the Invention**

This invention relates to a method of handling multiple backup processes. In one embodiment, the invention provides cascaded multiple target flashcopy consistency group start optimisation for background copy processes.

**2. Description of the Related Art**

The storage of data in large organisations is of fundamental importance, both for reliability of the data and for the ability to recover data in the event of any hardware failure. Storage area network (SAN) is an architecture that is used when very large amounts of data are needed to be stored in a reliable and secure manner. This technology allows networks to be created that support the attachment of remote computer storage devices such as disk arrays to servers in such a way that, to the operating system, the devices appear as locally attached. It is common in these networks to include a large amount of redundancy, both in the data storage and in the hardware connections between the individual components.

**SUMMARY OF THE INVENTION**

According to an aspect of the present invention, there is provided a method of handling multiple backup processes comprising creating a cascade of target storage volumes from a plurality of target storage volumes by inserting each most recent target storage volumes between a last added target storage volume and a source storage volume, the most recent target storage volumes each having a copy rate faster than the last added target storage volume, such that each added target storage volume has a faster copy rate than an immediately prior added target storage volume, and a most recent target storage volume added to the cascade of target storage volumes has a fastest copy rate; and receiving one or more instructions initiating a plurality of backup processes from the source storage volume to a plurality of target storage volumes.

In addition to the foregoing exemplary embodiment, various other system and computer program product embodiments are provided and supply related advantages. The foregoing summary has been provided to introduce a selection of concepts in a simplified form that are further described below in the Detailed Description. This Summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used as an aid in determining the scope of the claimed subject matter. The claimed subject matter is not limited to implementations that solve any or all disadvantages noted in the background.

**BRIEF DESCRIPTION OF THE DRAWINGS**

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

FIG. 1 is a schematic diagram of a pair of storage volumes;

FIG. 2 is a schematic diagram of a flashcopy cascade;

FIG. 3 is a schematic diagram of a plurality of storage volumes;

FIG. 4 is a schematic diagram of the construction of a cascade from the storage volumes of FIG. 3;

FIG. 5 is a schematic diagram of the ordering of the starting of backup processes; and

FIG. 6 is a flowchart of a method of handling the backup processes.

**DETAILED DESCRIPTION OF THE DRAWINGS**

Various methods exist for creating data redundancy. For example, a function such as a flashcopy function enables an administrator to make point-in-time, full volume copies of data, with the copies immediately available for read or write access. The flashcopy can be used with standard backup tools that are available in the environment to create backup copies on tape. A flashcopy function creates a copy of a source volume on a target volume. This copy, as mentioned above, is called a point-in-time copy. When a flashcopy operation is initiated, a relationship is created between a source volume and target volume. This relationship is a "mapping" of the source volume and the target volume. This mapping allows a point-in-time copy of that source volume to be copied to the associated target volume. The relationship exists between this volume pair from the time that the flashcopy operation is initiated until the storage unit copies all data from the source volume to the target volume, or the relationship is deleted.

When the data is physically copied, a background process copies tracks from the source volume to the target volume. The amount of time that it takes to complete the background copy depends on various criteria, such as the amount of data being copied, the number of background copy processes that are running and any other activities that are presently occurring. The flashcopy function works in that the data which is being copied does not actually need to be copied instantaneously, it only needs to be copied just prior to an update causing an overwrite of any old data on the source volume. So, as data changes on the source volume, the original data is copied to the target volume before being overwritten on the source volume.

Therefore, a flashcopy is a feature supported on various storage devices that allows a user or an automated process to make nearly instantaneous copies of entire logical volumes of data. A copy of a source disk is made on a target disk. The copies are immediately available for both read and write access. A common feature of flashcopy like implementations is the ability to reverse the copy. That is, to populate the source disk of a flashcopy map with the contents of the target disk. It is also possible to use flashcopy in cascaded implementations, in which a target disk later becomes the source disk for a further flashcopy or vice versa.

A cascaded configuration of storage volumes is described in detail in U.S. Pat. No. 7,386,695. It is also possible to create multiple cascades of storage volumes which are interlocking at a logical level. A first cascade may comprise storage volumes A, B, C and D which are arranged in a cascade as follows: A ↔ B ↔ C ↔ D, while at a later time a new backup of A may be started that ultimately leads to the creation of A ↔ E ↔ F. Many different combinations of flashcopy functions and reversed functions are possible, potentially creating complicated multiple cascading storage volumes.

A typical use case for multiple target flashcopy is to create a number of application consistent snapshots of a production disk (the source volume of the flashcopies). This is done using

some form of a consistency group in order to guarantee that all the copies are identical. Both the cascaded and the “classical” implementations of multiple target flashcopy do not minimise the read access redirected to the production disk effectively as part of the background copy process used to create the independent copies.

One solution to this problem, which is appropriate to both the classical and cascaded approaches to multiple flashcopy implementations would be to perform a single read of the source disk and multiple writes to each of the targets. However such a solution is not scalable with the number of targets created. This solution also makes it difficult to have different background copy rates for different flashcopy maps. Another solution appropriate to the cascaded approach is to ensure that the last started map completes its background copy process before the background process of the next last started map begins. The problem with this solution is that the length of time taken to produce an independent copy increases with the number of copies required.

Owing to the invention, it is possible to provide system that will schedule a cascade of backup processes (such as flashcopy) that will minimise the number of reads that are required of the host disk (the source storage volume), but without slowing down any of the actual backup processes. The solution improves the cascaded implementation of multiple target flashcopy by ordering the adding to the cascade of the flashcopy maps for a single source in a consistency group inversely with respect to the background copy rate of the map and suppressing the start of a flashcopy map’s background copy process until the background copy process of each map added to the cascade later has been started.

The advantage of this approach is that each map can have its own background copy rate, each map does not need to wait for “younger” maps to complete the background copy process and the solution is scalable with respect to the number of copies made. The total number of read and write operations required to complete the background processes for all maps is unchanged, but the number of read operations directed to the production disk is reduced.

For example, with a production disk A it is possible to make copies B, C, D and E of this disk using flashcopy, by creating flashcopy maps  $A \rightarrow B$ ,  $A \rightarrow C$ ,  $A \rightarrow D$  and  $A \rightarrow E$  in a single consistency group with background copy rates  $w$ ,  $x$ ,  $y$ ,  $z$  respectively. If the copy rates of the maps are such that  $w \geq x \geq y \geq z$ , then when the consistency group is started, the map  $A \rightarrow E$  is added to the cascade first, because this has the lowest background copy rate,  $z$  and so on. This results in cascade  $A \rightarrow B \rightarrow C \rightarrow D \rightarrow E$ . When all maps are in the cascade, the background copy process of  $A \rightarrow B$  is started first, followed by the backup process for  $A \rightarrow C$  and so on. This means that the backup process  $A \rightarrow B$  copies a grain of data from volume A to volume B first. Then backup process  $A \rightarrow C$  will copy the data from  $B \rightarrow C$  and so on. So it can be seen that only one read of the source storage volume (the production disk A) is required for all four copies to be completed, because the background copy rates of “younger” maps are greater than “older” ones.

FIG. 1 illustrates the concept of a flashcopy using a storage controller 8 and two storage disks 10 and 12. The disks 10 and 12 could form part of a larger array of disks, and would typically form part of an enterprise storage solution. The disks 10 and 12 could be part of a storage solution relating to a commercial website, for example. If at any time a backup needs to be made of the content of vdisk1, then a flashcopy instruction can be sent from the storage volume controller 8 to that disk 10, which defines a source disk 10 (vdisk1) and also a target disk 12 (vdisk2), which is the target of the flashcopy.

The flashcopy instruction creates a point-in-time copy of the image of the specific vdisk which is the source disk 10.

In the example of FIG. 1, the source disk 10 of a first flashcopy instruction is vdisk1, and the target disk 12 is vdisk2. The flashcopy instruction starts the flashcopy process, which creates a map 14 from the source disk 10 to the target disk 12. This map is labelled map 1 in the Figure. The image of vdisk1 at this specific point in time is now available on vdisk2. This creates a backup of the data on vdisk1, and also allows tests and other administration tasks to be run on the data of vdisk1, without the attendant danger of losing any of the original data, as it is preserved on the original source disk.

When a flashcopy is made, it creates a link between the two disks 10 and 12, as defined by the map 14. Data may now be copied across in the background, with the additional requirement that any access to vdisk2 (as the target disk 12) will immediately cause the relevant parts of the image of vdisk1 to be copied across, and also any access to vdisk1 which would result in a change to the image stored by that disk 10 will also cause the unaltered data to be immediately copied across to the target disk 12, prior to the change being made. In this way, the vdisk2, to an outside user, stores the point in time copy of vdisk1, although data will only be physically copied across under the circumstances described above.

A storage volume that is the target volume of a flashcopy function can also be the source volume of a further flashcopy function, thus creating a cascade of storage volumes. In FIG. 2 there is shown an example of a flashcopy cascade of three storage volumes 10, 12 and 16, which are linked by respective flashcopy maps 14. Each map 14 defines a flashcopy function from a source volume to a target volume. Disk B is providing a backup of disk A, and disk C is providing a backup of disk A, through disk B. The flashcopy functions 14 linking the different storage volumes may have been started at different times, which create different point-in-time copies of the images stored by the respective storage volumes, or could have been started simultaneously.

In the flashcopy cascade of  $A \rightarrow B \rightarrow C$ , where A, B and C are the disks in the cascade, as shown in FIG. 2, and the arrows are the flashcopy maps, then denoting  $(A, B)$  to be a flashcopy mapping from disk A to disk B, the cascade has maps  $(A, B)$  and  $(B, C)$ . In this implementation of the cascade, any new data write to disk A will cause a write to disk B, as per the respective flashcopy function, which is required to maintain the image on disk B. This writing to disk B this will cause a further clean read of disk B followed by a write to disk C. In this way a single write to the first storage volume 10 in the cascade can result in a number of operations throughout the cascade.

When a cascade is created, the new maps and new storage volumes are inserted into the cascade, not added to the end of the cascade. In the cascade shown in FIG. 2, the first backup process started would be  $A \rightarrow C$ . When the backup process  $A \rightarrow B$  is then started, the new target storage volume B is effectively “inserted” between the existing source storage volume A and the existing target storage volume C. This “insertion” is purely a logical construction illustrating the fact that target disk C will receive data writes from disk B, rather than disk A. This is how a cascaded implementation differs from a conventional arrangement which would have two independent maps from disk A.

When a cascade is formed by the simultaneous taking of multiple backups from the disk A, then an ordering of the cascade is created using the different copy rates of the background backup processes that have been started. For example, a user may simultaneously request that three different copies

5

be made of a production disk A. The different copies may be taken for different purposes. A first backup may be a true backup of the data on disk A in order to guard against any failure of disk A. A second backup may be taken to be used to test certain processes on the data, without the danger of losing any of the original data, while a third backup may be taken for the purpose of reading the data to tape, to create a permanent record.

The different backup processes can therefore have different copy rates associated with them, which can be user defined, or more likely will be initially set by administrator at a system level and used automatically according to the purpose of the backup process being taken. For example, the backup process initiated to perform testing on the stored data will have a higher copy rate assigned, as the backup of the data on disk A, for this purpose, needs to be completed as quickly as possible, in order to allow the testing to start as soon as possible. Contrastingly, the backup process simply to make a copy of the data on a different disk as a true backup of the data may have a lower copy rate assigned.

FIG. 3 illustrates a plurality of storage volumes 10 to 18 to be used to create a flashcopy cascade that is to be ordered with respect to three different backup processes, which can be characterised by different copy rates “fast”, “slow” and “medium”. These copy rates are relative terms and define the rate at which data is copied from the source storage volume A (the actual production disk) to the individual target storage volumes. Although the copy rates are shown as labelling the specific storage volumes, it should be understood that the copy rates relate to the backup process that is copying the data to the specific storage volume, not the actual storage volume itself. In hardware terms, the target storage volumes are identical.

The process of creating the cascade from the storage volumes of FIG. 3 is shown in FIG. 4. Each target storage volume is added in turn to the cascade of target storage volumes that cascades from the source storage volume, in such a way that the target storage volumes are added to the cascade in an order that is inversely proportional to the copy rate of the respective backup process. It can be seen in FIG. 4, that the first action taken is the adding of target storage volume D is first added to the cascade, as this storage volume has the slowest copy rate of the backup process that is associated with that specific storage volume.

The second action is the addition of the target storage volume C, which has “medium” copy rate for the backup process that is associated with that specific storage volume C. The cascade now contains three storage volumes, the original production disk A (which is the source storage volume) and the two target storage volumes C and D. Note that the nature of the cascade addition rule is that the last added storage volume is placed closest to the source storage volume A, so that the addition of the target storage volume C means that this disk is inserted between volumes A and D, in the logical arrangement of the storage volumes.

Finally, the third target storage volume, disk B, is added to the cascade. This storage volume has the fastest copy rate for the backup process that is associated with the specific disk B, so this target storage volume is added to the cascade last, as the volumes are added in an order that is inversely proportional to their associated copy rate. As per the algorithm for adding disks to a cascade, the last added storage volume, disk B, is placed closest to the source storage volume A. The storage volume B is therefore inserted between the production disk A and the target storage volume C. This action completes the cascade.

6

Once the cascade has been created, then the backup processes that define the maps between the storage volumes need to be started. The storage volume controller 8 will start each backup process in turn, the backup processes being started in an order from the most recent target storage volume added to cascade to the first target storage volume added to cascade. As can be seen in FIG. 5, this means that the backup process that is started first is the process that is associated with the target storage volume B, as this was the last storage volume to be added to the cascade. As can be seen in the Figure, this backup process is started first.

The remaining backup processes from the production disk A, which form the cascade, are then started in turn. After the backup process for storage volume B has been started then the backup process for volume C is started, followed by the start of the backup process for volume D. In this way, the backup processes are started one after another until all have been started. The order in which the processes are started is the opposite of the order by which the different storage volumes were added to the cascade. Faster backup processes are started before slower backup processes. Once the cascade has been ordered and the backup processes started, then the flash-copy cascade will run as normal.

The methodology described above for the addition of the volumes to the cascade and the starting of the backup processes must also be able to deal with the situation when two (or more) backup processes that are scheduled to start together actually have the same copy rate. This can easily happen if, for example, a user wishes to perform two different testing scenarios on the dataset of production disk A and will correspondingly schedule two backup processes from disk A to respective target storage volumes. The storage volume controller 8 can make an arbitrary decision on which to add to the cascade first, as this is immaterial to the overall performance, but must still start the backup processes in the inverse order to which the volumes were added to the cascade.

FIG. 6 is a flowchart summarizing the operation of the storage volume controller 8 in performing the flashcopy scheduling function. The first step S1 comprises receiving one or more instructions initiating a plurality of backup processes from a single source storage volume to a plurality of target storage volumes. All of the backup processes may be triggered by a single instruction from a user, or multiple instructions may be sent by the user to initiate the taking of the backups from the production disk. Essentially multiple copies are to be taken from the same source storage volume with multiple independent backup processes. The storage volume controller 8 is responsible for the scheduling of these tasks.

The second step S2 performed by the storage volume controller 8 comprises the step of adding each target storage volume to a cascade of target storage volumes from the source storage volume, the target storage volumes being added to the cascade in an order that is inversely proportional to the copy rate of the respective backup process. This step is described above in detail with reference to the example of FIG. 4. Essentially the target storage volumes are added in order to the cascade, with the volume associated with the slowest copy rate being added to the cascade. After each individual addition, then the next volume associated with the next slowest copy rate is added to the cascade.

The final step of the method is step S3 in which each backup process is started in turn, the backup processes being started in an order from the most recent target storage volume added to cascade to the first target storage volume added to cascade. In a preferred embodiment, this step of starting each backup process in turn comprises suppressing the start of each backup process until all of the backup processes for

target storage volumes added later to the cascade have been started. This step is described above in detail with reference to the example of FIG. 5. The backup processes are each started in turn, in the opposite order to which the corresponding storage volumes were added to the cascade.

The scheduling of the cascade of backup processes described above minimizes the number of reads that are required of the production disk, but without slowing down any of the actual backup processes. This solution provides an improvement in the cascaded implementation of multiple target flashcopy. This is achieved by ordering the adding to the cascade of the flashcopy maps for a single source in a consistency group inversely with respect to the background copy rate of the map and suppressing the start of a flashcopy map's background copy process until the background copy process of each map added to the cascade later has been started. The main advantage of this approach is that each background process will have its own background copy rate and does not need to wait for other background copy processes to complete. The total number of read and write operations required to complete the background processes for all backup processes is unchanged, but the number of read operations directed to the production disk is reduced.

While at least one exemplary embodiment has been presented in the foregoing detailed description of the invention, it should be appreciated that a vast number of variations exist. It should also be appreciated that the exemplary embodiment or exemplary embodiments are only examples, and are not intended to limit the scope, applicability, or configuration of the invention in any way. Rather, the foregoing detailed description will provide those skilled in the art with a convenient road map for implementing an exemplary embodiment of the invention, it being understood that various changes may be made in the function and arrangement of elements described in an exemplary embodiment without departing from the scope of the invention as set forth in the appended claims and their legal equivalents.

The present invention may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable program instructions for storage in a computer readable storage medium within the respective computing/processing device.

Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code written in any combination of one or more programming languages, including an object oriented programming language such as Smalltalk, C++ or the like, and conventional procedural programming languages, such as the "C" programming language or similar programming languages. The computer readable program instructions may execute entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

These computer readable program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

The computer readable program instructions may also be loaded onto a computer, other programmable data processing

apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). In some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts or carry out combinations of special purpose hardware and computer instructions.

While one or more embodiments of the present invention have been illustrated in detail, the skilled artisan will appreciate that modifications and adaptations to those embodiments may be made without departing from the scope of the present invention as set forth in the following claims.

The invention claimed is:

1. A method of handling multiple backup processes comprising:

creating a cascade of target storage volumes from a plurality of target storage volumes by inserting each most recent target storage volumes between a last added target storage volume and a source storage volume, the most recent target storage volumes each having a copy rate faster than the last added target storage volume, such that each added target storage volume has a faster copy rate than an immediately prior added target storage volume, and a most recent target storage volume added to the cascade of target storage volumes has a fastest copy rate; and

receiving one or more instructions initiating a plurality of backup processes from the source storage volume to a plurality of target storage volumes.

2. The method of claim 1, further including starting each backup process in turn, the backup processes started in an order from a most recent target storage volume added to the cascade of target storage volumes to the first target storage volume added to the cascade of target storage volumes.

3. A method according to claim 2, wherein the step of starting each backup process in turn comprises suppressing a start of each backup process until all backup processes for target storage volumes subsequently added to the cascade have been started.

4. A method according to claim 3, further comprising detecting that two or more backup processes have a same copy rate and, when adding each target storage volume to the cascade of target storage volumes, making an arbitrary decision on which of the target storage volumes with backup processes having the same copy rate should be added to the cascade of target storage volumes first.

5. A method according to claim 2, wherein the step of starting each backup process in turn comprises suppressing a start of each backup process until all backup processes for target storage volumes subsequently added to the cascade have been started.

6. A method according to claim 1, wherein each backup process comprises a flashcopy function creating a point-in-time copy of a first source storage volume on a corresponding target storage volume.

7. A method according to claim 6, further comprising detecting that two or more backup processes have a same copy rate and, when adding each target storage volume to the cascade of target storage volumes, making an arbitrary decision on which of the target storage volumes with backup processes having the same copy rate should be added to the cascade of target storage volumes first.

8. A method according to claim 1, further comprising detecting that two or more backup processes have a same copy rate and, when adding each target storage volume to the cascade of target storage volumes, making an arbitrary decision on which of the target storage volumes with backup processes having the same copy rate should be added to the cascade of target storage volumes first.

\* \* \* \* \*